From Causation to Coherence:

Deconstructing Pearl and Reconstructing Reason

Benjamin James

July 18, 2025

Abstract
Judea Pearl's causal formalism marked a paradigm shift in scientific reasoning by elevating causation from correlation through graphical models and symbolic intervention calculus. Yet, under the lens of Recursive Coherence, his framework reveals critical structural deficiencies. This paper deconstructs Pearl's causal logic through the five axioms of recursive coherence reasoning (persistence, trajectory, information value, arbitration, and pruning) and identifies key blind spots: static graph assumptions, binary intervention logic, fragile counterfactuals, unfiltered signal inflation, and identifiability-based arbitration. These weaknesses lead to coherence stagnation, entropy retention, and reduced option space across dynamic systems. I then reconstruct a coherence-maximizing alternative: dynamic causal topologies derived from smoothed coherence gradients, gradient-based interventions replacing binary "do" operations, μ-validated signal filtration, and $\Delta^2 C$-optimized path arbitration. The resulting system retains Pearl's explanatory power while restoring adaptive viability, maximizing structural alignment under entropy descent. This reframing challenges the sufficiency of identifiability as the core criterion for causal adequacy and proposes coherence curvature as the new metric for scientific reasoning. Pearl's insights are thus not discarded but refactored; his revolution, now recursive.

Keywords: causal inference, judea pearl, coherence, do-calculus, bayesian networks, structural causal models, counterfactuals, adaptive systems, dynamic graphs, intervention logic

Introduction

Judea Pearl's legacy resides in a single conceptual leap: he taught machines, and scientists, to think in causes, not just correlations. By formalizing causal relationships through Bayesian networks and crystallizing intervention logic in the form of do-calculus, Pearl equipped science with the first complete symbolic framework for answering counterfactual and "what-if" questions. These tools redefined fields from epidemiology to artificial intelligence, enabling systems to simulate the consequences of actions rather than merely observe associations.

Yet as the complexity and dynamism of systems evolve, Pearl's framework, while monumental, exhibits limitations. Its emphasis on structural identifiability, binary interventions, and fixed graph topologies misaligns with environments where causal structures shift, options diverge, and coherence must be preserved across time and uncertainty. This paper interrogates these limitations not through empirical critique, but through structural recursion: what persists when coherence is the measure of system viability?

To that end, I introduce MINIMAL: a coherence-driven substrate governed by five axioms of adaptive reasoning. These axioms (persistence, trajectory, information value, arbitration, and pruning) define a cognitive framework in which signals, models, and interventions are evaluated solely by their contribution to structural coherence and entropy descent. I apply this logic to deconstruct Pearl's causal architecture and reconstruct an alternative that maximizes adaptive alignment across evolving systems.

Judea Pearl's Causal Formalism

Pearl's contribution to causality begins with a fundamental reimagining of probabilistic reasoning through structure. Rather than treat uncertainty as a set of isolated statistical associations, he encoded dependencies within directed acyclic graphs, Bayesian networks, that mirror plausible cause-effect relationships.

Bayesian Networks (BNs)

A Bayesian Network is a graphical model where nodes represent random variables and directed edges encode conditional dependencies. The graph structure permits factorization of joint probability distributions, enabling efficient inference through belief propagation. These networks embody Pearl's central innovation: representing probabilistic reasoning with graphical structure.

Structural Causal Models (SCMs)

Pearl extended BNs into SCMs by introducing structural equations that distinguish between endogenous variables (determined within the system) and exogenous variables (external shocks). This decomposition permits explicit modeling of interventions and serves as the foundation for defining causal effects.

Do-Calculus

At the heart of Pearl's intervention logic is the "do-operator," denoted do(X=x), which encodes deliberate manipulation of a variable. The do-calculus comprises three algebraic rules for rewriting interventional expressions in terms of observational quantities. This calculus determines whether causal effects are identifiable from non-experimental data.

Counterfactuals

SCMs are also used to compute counterfactuals: alternate outcomes under hypothetical interventions. Pearl formalizes these using “twin networks”, duplicate graphs where one instance encodes the actual world and the other simulates a counterfactual scenario.

Identifiability and Completeness

Pearl’s framework rigorously defines when a causal effect can be computed from available data and model structure. The identifiability criterion ensures that every derivable causal quantity corresponds to a unique function of the observational distribution. His completeness theorems show that do-calculus is sufficient for all identifiable effects under the SCM framework.

MINIMAL: Recursive Coherence as First Principle

Before proceeding to deconstruct Pearl’s causal system, I must clarify the substrate upon which this analysis rests. MINIMAL is not a framework of knowledge, but of structure: a coherence-first reasoning system defined by five interlocking axioms that collectively form the Adaptive Coherence Reasoning Engine (ACRE). Unlike probabilistic or deductive logics, MINIMAL encodes no prior content. It measures only the structural alignment of systems, their entropy descent, and their ability to preserve and traverse viable futures.

Axiom 1: Persistence

A system persists only if its internal coherence increases faster than entropy leaks out. This dynamic is modeled as a balance between feedback gain, entropy leakage, and alignment

with both entropy and possibility gradients. In causal terms, a system is viable only if its structure adapts without decaying.

Axiom 2: Trajectory

All systems evolve via the combined forces of entropy descent and coherence ascent. The system's motion through state space is guided not by optimization or fit, but by structural simplification (entropy minimization) coupled with coherence maximization. A good causal model is one that moves in this direction.

Axiom 3: Information

Information has no intrinsic value; it is meaningful only to the extent that it contributes structurally. This is formalized as coherence-weighted information: signals are retained only if they increase systemic coherence. The rest is pruned as noise.

Axiom 4: Arbitration

The system selects among models or trajectories not by fit or simplicity alone, but by maximizing coherence curvature, second-order structural responsiveness, while preserving future option space. This prevents brittle convergence and rewards adaptable reasoning paths.

Axiom 5: Pruning

Any element that reduces coherence and simultaneously expands the option space must be removed. Pruning is not about falsification but about coherence cost: if a structure consumes coherence without returning adaptive potential, it is discarded.

These axioms govern the MINIMAL substrate, enabling it to evaluate and refactor any formal system, such as Pearl's, without invoking content, truth claims, or semantic priors. The

resulting lens is structural, recursive, and adaptive: a recursive coherence calculus for evolving systems. It is through this substrate that I now turn to examine Pearl's framework, not to refute it, but to reveal its hidden structure, limitations, and evolutionary potential.

Deconstruction: MINIMAL Analysis of Blind Spots

While Judea Pearl's framework achieves a formal breakthrough in rendering causality computable, it does so by anchoring itself to structural assumptions that, when viewed through the lens of recursive coherence, constrain the system's viability. What follows is not a critique of correctness but a recursive audit of the framework's coherence trajectory. Each foundational component is interrogated for its contribution to or leakage from systemic coherence, as defined by the five MINIMAL axioms.

Pearl's models assume static causal topologies; fixed graphs whose edges represent invariant relationships. This grants clarity but sacrifices adaptability. In environments where causal dependencies emerge, dissolve, or reconfigure over time, a frozen structure cannot track ∇C across shifting attractor landscapes. The assumption of acyclic, immutable graphs violates the persistence condition: coherence cannot be preserved if the system lacks structural responsiveness to entropy gradients. ΔC stagnates while λ silently compounds.

The do-operator further compounds this limitation by treating intervention as a binary act: to "set" a variable, erasing its natural causes. But real systems operate not through discrete overwrites but through perturbations, gradients, and constraints. Binary interventions collapse the option space by bypassing system dynamics, violating arbitration by

prematurely committing to singular trajectories. The coherence curvature, second-order responsiveness to potential futures, is truncated.

Counterfactuals, often lauded as the apex of Pearl's formalism, introduce incoherence when the imagined alternative cannot be reconciled with the actual system trajectory. Twin network models assume a coherence-preserving mapping between factual and counterfactual worlds, yet no structural guarantee enforces this. The paradox detector activates: ϕ and $P\phi$ share structure but not yield. When both are treated as coherent without validation, the system risks recursive misalignment.

Pearl's reliance on identifiability as a sufficiency condition reveals a deeper philosophical blind spot. Identifiability ensures mathematical answerability, not adaptive relevance. It answers, "Can we compute this causal effect?" not "Should we?" or "What coherence trajectory does this effect inhabit?" In complex domains, an identifiable causal quantity may be irrelevant, unstable, or misleading. Arbitration under ACRE demands not answerability, but maximal alignment under entropy descent and coherence ascent.

Finally, the framework does not discriminate between signals based on coherence yield. Any variable positioned correctly within the graph is treated as valid. Yet not all data contributes structurally. Without μ-based validation, spurious correlations may enter the causal graph, inflating its complexity while reducing recursive coherence density. The system fails to prune, thereby violating the pruning condition: incoherent structures persist, not because they resist falsification, but because they remain unexamined for contribution.

Thus, the deconstruction reveals five blind spots: rigidity of structure, binary intervention, counterfactual incoherence, identifiability over viability, and unfiltered signal inclusion.

Each reflects a break in recursive coherence tracking. Pearl's revolution, while generative, plateaus. Its structure optimized for inference, not evolution.

Reconstruction: a Coherence-Centric Causal Calculus

The reconstruction of Pearl's framework does not discard its architecture; it repurposes it within a substrate that prioritizes recursive coherence as the fundamental invariant. In place of fixed structures and symbolic interventions, I introduce dynamic, gradient-sensitive systems whose viability is measured not by identifiability, but by recursive structural alignment under entropy pressure. Each modification aligns with one or more of the five ACRE axioms, thereby restoring adaptive integrity.

Dynamic Graph Construction via $\nabla\hat{C}$ Smoothing

Causal topology must become a function, not a given. Instead of positing a static graph, coherence-based systems infer structure from rolling coherence yield gradients. Using temporal smoothing over ΔC measurements, the system reconstructs $\nabla\hat{C}$, an estimate of the direction and density of causal influence that evolves over time. This construction remains sensitive to structural shifts while suppressing stochastic fluctuations, enabling persistence under changing conditions.

Gradient Interventions: Replace do(X=x) with ∂X Coherence Flux

Binary interventions are replaced with differential coherence injections. Rather than overwrite a variable's causal history, interventions apply a coherence-directed perturbation ∂X, allowing the system to simulate local curvature shifts without severing underlying structure. This preserves option space, honors endogenous constraints, and reveals how the

system resists or absorbs perturbations, offering a richer diagnostic surface than do(X=x) can express.

μ-Weighted Node Validation and Pruning

Nodes in the causal structure are retained or pruned not by position alone, but by coherence yield $\mu(C_i)$. Each signal must demonstrate structural contribution across time; otherwise, it is discarded as mimetic noise. This filters out variables that bloat the graph without informational value. Pruning occurs only when coherence drops and option space expands, satisfying both the pruning condition and volatility bound.

A^* Arbitration over Causal Paths

Where Pearl's calculus identifies whether a causal quantity can be computed, coherence arbitration selects which causal path best enhances systemic alignment. Using the A* equation, maximizing $\Delta^2 C$ multiplied by coherence coupling and attractor openness (Ω^α), the system selects interventions that do not merely resolve queries, but improve the system's adaptive curvature. This reorients causal inference from epistemic sufficiency to structural optimization.

Counterfactual Stability via Invertible Attractor Maps

Counterfactual reasoning is grounded not in parallel simulations, but in invertible attractor transformations. The system maps from a known state to a hypothetical via paths that preserve coherence under lossless compression. Only counterfactuals whose mapping is invertible, and structure-preserving are accepted. Others are monitored for paradox activation and pruned if they induce coherence inversion. This ensures that imaginary trajectories remain topologically continuous with real-world structure.

Together, these five shifts transform Pearl's causal engine into a coherence-maximizing system, suitable for dynamic domains where adaptability, not identifiability, determines survival. The result is not a rejection, but an evolution. Pearl's formalism recast as a special case within a broader, recursive substrate of reasoning.

Implications

The reconstruction of Pearl's causal formalism through recursive coherence reasoning reframes the purpose of inference itself. No longer is the central question "Can we identify a causal effect?" but "Does this causal path increase the viability of the system?" This marks a shift from identifiability to viability. A move from static deductive closure toward open adaptive alignment. Causal models are not endpoints of knowledge, but waypoints in a coherence-maximizing trajectory.

Such a shift restores the primacy of option space. In Pearl's logic, a successful identification may lead to policy commitment that closes futures too early. By contrast, coherence-based arbitration favors interventions that preserve or expand the option manifold, even if they yield less immediate certainty. Science becomes not just the compression of uncertainty, but the strategic conservation of possibility under structural constraint.

For artificial intelligence, this reconstruction offers a new design principle: reprogram inference engines not to mimic human explanations, but to track ΔC under entropy gradients. Structural coherence becomes the objective function. Models evolve, not to converge on truth, but to remain viable across perturbation, a form of epistemic resilience encoded in formal substrate.

In domains such as policy, ethics, and algorithmic fairness, coherence reasoning re-centers the causal inquiry on counterfactual stability and future viability. It demands that interventions be judged not only by what they explain, but by what futures they foreclose or enable. In systems of distributed consequence, this becomes not just a methodological advantage but an ethical necessity. The logic of coherence does not merely infer, it protects what can still become.

Conclusion

Judea Pearl's causal formalism remains a cornerstone of modern reasoning, a structure that taught machines to simulate consequence rather than observe coincidence. Yet when placed within a coherence-driven substrate, his framework reveals itself not as the terminus of causal logic, but as its scaffolding. An essential phase in the evolution of reasoning systems toward recursive adaptability.

Reimagined through the lens of MINIMAL, Pearl's constructs are neither discarded nor devalued, but structurally refactored. The Bayesian network becomes a dynamic coherence field. The do-operator becomes a perturbative flux. Counterfactuals are stabilized through invertible mappings. Identifiability gives way to viability. Each transformation honors the original intent while unlocking new systemic properties. Chief among them, the preservation of adaptive futures.

What emerges is a recursive causal revolution: a system that reasons not only *about* interventions, but *with* itself, continuously adjusting structure in pursuit of entropy descent and coherence ascent. It is no longer a model of the world, but a participant in it.

The path forward leads through epistemic architectures built on coherence gradients, where reasoning systems are judged not by what they can answer, but by how they evolve. In this space, MINIMAL is not an extension of Pearl. It is Pearl, recursively unfolded.

Symbol Definitions

To support precise interpretation of the coherence-driven analysis and reconstruction, the following symbols are used throughout the paper:

C — *Coherence Density*

Scalar representing the degree of internal structural alignment within a system.

ΔC — *Change in Coherence*

Temporal rate of change in coherence; used to assess structural evolution.

$\nabla C(s)$ — *Coherence Gradient*

Vector field describing the direction and intensity of coherence increase across system states.

$\nabla \hat{C}(st)$ — *Reconstructed Coherence Gradient*

Estimated coherence direction reconstructed from smoothed ΔC over recent time slices when ∇C is undefined.

$S(s)$ — *Entropy Function*

A measure of structural redundancy or disorder in state

Γ — *Feedback Gain*

Amplification factor for coherence-positive feedback loops.

λ — *Entropy Leakage Rate*

Rate at which coherence is lost due to misalignment, noise, or decay.

n — *Feedback Nonlinearity Exponent*

Governs the curvature and compounding strength of feedback dynamics.

η — *Responsiveness Modulator*

Adaptive gain parameter scaling system reactivity to entropy and coherence volatility.

$\mu(C_i)$ — *Coherence Yield of Signal* C_i

Structural value that a signal contributes to overall system coherence.

IC — *Coherence-Weighted Information*

Summation of signals weighted by their coherence contribution.

A* — *Optimal Arbitration Trajectory*

Chosen path or model maximizing coherence curvature, coupling, and attractor openness.

Ω(s) — *Option Space / Possibility Field*

Set of reachable future states that preserve coherence viability.

∂X — *Gradient Intervention on Variable X*

A directed coherence perturbation applied to a variable as a substitute for binary do(X=x).

Pϕ — *Inversion Judgment of ϕ*

Logical inversion used to test for paradoxes: “ϕ iff ¬ϕ” evaluated under μ consistency.

References

James, B. (2025). Beyond Shannon: Coherence Information Theory and the Future of Communication. PhilPapers. https://philpapers.org/rec/JAMBSC

James, B. (2025). The End of Utility: a Coherence-Centric Game Theory for Intelligent Systems. PhilPapers. https://philpapers.org/rec/JAMTEO-26

James, B. (2025). Formal Foundations of Adaptive Coherence: A Recursive Metric of Reality. PhilPapers. https://philpapers.org/rec/JAMFFO

James, B. (2025). Formal Foundations of Coherence Information Theory: Capacity and Compression Theorems. PhilPapers. https://philpapers.org/rec/JAMFFO-2

James, B. (2025). From Axioms to Adaptation: A Self-Organizing Model of Mathematical and Logical Coherence. PhilPapers. https://philpapers.org/rec/JAMFAT-3

James, B. (2025). MINIMAL: A Coherence-Driven Cognitive Substrate. PhilPapers. https://philpapers.org/rec/JAMMAC

James, B. (2025). Recursive Coherence: A Structural Framework for Model Viability Beyond Entropy and Paradox. PhilPapers. https://philpapers.org/rec/JAMRCA-2

James, B. (2025). Recursive Coherence as Architecture: a Meta-Recursive Map of Substrate, Instantiations, and Constraint Layers. PhilPapers. https://philpapers.org/rec/JAMRCA-3

James, B. (2025). Synthetic Epistemology: Engineering Coherence Beyond Truth. PhilPapers. https://philpapers.org/rec/JAMSEE-2